

REPORT DOCUMENTATION PAGE

Form Approved
OMB NO. 0704-0188

Public Reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comment regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE July 2001	3. REPORT TYPE AND DATES COVERED Final Report, 03/02/98 - 03/02/01
4. TITLE AND SUBTITLE Multithreaded User Space Multiprotocols: Experiments with Collaborative Multimedia		5. FUNDING NUMBERS DAAG55-98-1-0125 ARO CI 38245	
6. AUTHOR(S) V. Rego			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Purdue University, West Lafayette, IN 47906		8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U. S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211		10. SPONSORING / MONITORING AGENCY REPORT NUMBER 38425.1-CI-RIP	
11. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.			
12 a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution unlimited.		12 b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This research is motivated by emerging virtual collaboratories, distributed multimedia, and scalable distributed computing applications on modern heterogeneous networks of machines. It has had a humble origin, namely, a simple threads system (Ariadne) that was developed in our software laboratory (PacsLab) at Purdue. The Ariadne project has since expanded in scope, matured, and won recognition as a powerful supporting infrastructure for timer-enhanced multithreaded distributed computing. This is mainly due to its portability, efficiency and support for thread-migration. Ariadne has been used to support optimistic distributed simulations on large Sun and SGI workstation clusters, the IBM SP/2, and the Intel Paragon. It has led to the development of a novel sister threads system, (Arachne) which supports thread-migration on heterogeneous networked platforms. The funds provided by this award have enabled us to set up a multiprocessor cluster on which we can conduct experiments on thread migration, protocols and collaborative computing. These experiments are currently in progress.			
14. SUBJECT TERMS Threads, migration, protocols, collaborative computing.		15. NUMBER OF PAGES 2	
		16. PRICE CODE	
17. SECURITY CLASSIFICATION OR REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION ON THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL

NSN 7540-01-280-5500

Standard Form 298 (Rev.2-89)
Prescribed by ANSI Std. Z39-18
298-102

20011113 143

NOV 07 2001
gm

MASTER COPY: PLEASE KEEP THIS "MEMORANDUM OF TRANSMITTAL" BLANK FOR REPRODUCTION PURPOSES. WHEN REPORTS ARE GENERATED UNDER THE ARO SPONSORSHIP, FORWARD A COMPLETED COPY OF THIS FORM WITH EACH REPORT SHIPMENT TO THE ARO. THIS WILL ASSURE PROPER IDENTIFICATION. NOT TO BE USED FOR INTERIM PROGRESS REPORTS; SEE PAGE 2 FOR INTERIM PROGRESS REPORT INSTRUCTIONS.

MEMORANDUM OF TRANSMITTAL

U.S. Army Research Office
ATTN: AMSRL-RO-BI (TR)
P.O. Box 12211
Research Triangle Park, NC 27709-2211

☐ Reprint (Orig + 2 copies)

☐ Technical Report (Orig + 2 copies)

☐ Manuscript (1 copy)

☒ Final Progress Report (Orig + 2 copies)

☐ Related Materials, Abstracts, Theses (1 copy)

CONTRACT/GRANT NUMBER: DAAG55-98-1-0125

REPORT TITLE: Multithreaded User Space Multiprotocols: Experiments with Collaborative
Multimedia

is forwarded for your information.

SUBMITTED FOR PUBLICATION TO (applicable only if report is manuscript):

Sincerely,

Vernon Rego

REPORT DOCUMENTATION PAGE (SF298)
(Continuation Sheet)

In the initial stages of our work with threads, we found that layering the Ariadne system on the message-passing libraries PVM, P4 and LAM-MPI gave valuable functionality, but inadequate performance. We traced the runtime performance hits first to interactions at the threads-messaging interface, and later to protocol performance in the messaging subsystem itself. These performance problems are distinct from, and in addition to problems of scalability, system startup and functional inadequacy (e.g., efficient support for multiple media or realtime interactions) in extant message-passing systems like PVM, P4 and LAM-MPI. Thus, the need for improved performance and functionality motivated us to look for ways to make an application's messaging subsystem interact more efficiently with its computation subsystem.

The search brought us to user-space protocols, and to the conclusion that the cheapest and most efficient form of computation and communication interaction is via transparent user-space threads support for protocol actions. Contrast this approach to other systems (e.g., Nexus, TPVM, UPVM) which simply layer threads and threads-services upon extant messaging subsystems. Protocol layers, provided mainly for modularity, make protocol actions less sensitive to client needs. Layer boundaries tend to restrict interfaces, and lower layers are unable to query upper layers so as to make the best use of asynchronous actions, timing, and buffering. The frequent address-space crossings that are required by most in-kernel protocol implementations prevent applications from taking complete advantage of low-latency networks.

The funding for this project enabled us to purchase a high-end Sun enterprise multiprocessor server, approximately a year ago. We have since been experimenting with multiple processors, threads, timers, protocols and distributed computing --- bringing us to many interesting problems. We have been conducting experimental research in concepts and methods to provide portable and efficient threads-based multiprotocol support in user-space. This support is ideal for distributed multimedia and collaborative distributed applications, and is scalable in two senses --- in terms of nodes (via transaction-oriented connectionless transport), and in terms of multiprotocols (via integration of application-threads and multiple protocol-threads in user-space). The ideas are based on timer-enhanced threads for integration of applications and protocols, for scalable real-time and transaction-oriented communication. Further, we currently seek efficient ways of providing multiway communication for "group" interactions in collaborative distributed computing environments.

Based on a skeletal communications architecture called (CLAM), we have already made some headway. Early encouraging results include scalability characteristics superior to messaging systems based on TCP/IP, throughput and response times that are competitive with TCP/IP, a 100% improvement on PVM-UDP and LAM-UDP, and up to eight times the multicast throughput of PVM, all on Ethernets. Further, all this transparent communication-threads support is in addition to the benefits of threads for application development. We are currently working on ways to improve these early results on performance and functionality.

Publications

1. *Emulating Parallel Programming Environments in the Harness Metacomputing System*, (preprint) M. Migliardi and V. Sunderam.
2. *Automatic Reincarnation of deceased plug-ins in the Harness Metacomputing System*, (preprint) M. Migliardi and V. Sunderam.
3. *A Repository System with secure file access for collaborative environments*, (preprint) P. Gray, S. Chandramohan and V. Sunderam.
4. *On Tailoring Thread Schedules in Protocol Design: Experimental Results*, (preprint), J. Gomez, V. Rego and V. Sunderam.
5. *Multithreading Techniques and Applications*, (in prepration), J. Sang and V. Rego.